arXiv:1411.7565v2 [math.ST] 26 Jul 2016

Exact testing with random permutations

Jesse Hemerik and Jelle Goeman
Leiden University Medical Center, The Netherlands

July 27, 2016

Abstract

Permutation tests and other methods based on the permutation principle are highly popular, for instance in omics data analysis. Often random permutations are used, to limit the computation time. However, the existing theory on permutation tests usually assumes that all permutations are enumerated; little theory exists on testing with random permutations. It is known that naively using random permutations instead of the full permutation group, leads to anti-conservativeness. Also it has been claimed that adding the original observation to the random permutations solves this problem. A proof however is lacking. In this paper existing claims on validity of tests with random permutations are proven. Further, novel exact tests are presented.

Keywords and phrases: Permutation test, Nonparametric test.

1 Introduction

Permutation tests are nonparametric tests that can be used in particular when under the null hypothesis the data distribution is invariant under certain transformations (Fisher, 1936; Lehmann and Romano, 2005; Ernst et al., 2004). Apart from permutations, other groups of transformations can be used, such as rotations (Langsrud, 2005). Permutation tests are frequently used in for example omics and neuroscience.

Address: Department of Medical Statistics and Bioinformatics, Leiden University Medical Center, Postzone S5-P, Postbus 9600, 2300RC Leiden, The Netherlands.
E-mail: j.b.a.hemerik@lumc.nl

When the set of transformations used is not a group, a permutation test can be very conservative or anti-conservative. The first author who explicitly assumed a group structure is Hoeffding (1952). The role of the group structure has recently been emphasized (Southworth et al., 2009; Goeman and Solari, 2010). Southworth et al. (2009) note that in particular the set of 'balanced permutations' cannot be used, since it is not a group.

Often it is computationally infeasible to use the whole group of all permutations, due to its large cardinality. Hence in practice often random permutations are used, as was first proposed by Dwass (1957). Often a permutation p-value based on random permutations is simply seen as an estimate of the p-value based on all permutations.

It is known that naively using random permutations instead of all possible permutations can lead to extreme anti-conservativeness (Phipson and Smyth, 2010), especially when combined with multiple testing procedures. Therefore sometimes the identity permutation, which corresponds to the original observation, is included with the random permutations (Ge et al., 2003; Lehmann and Romano, 2005; Phipson and Smyth, 2010). Often this goes without explanation. Lehmann and Romano (2005) (p.636) state that when the identity is added, the estimated p-value is stochastically larger than the uniform distribution on $[0, 1]$ under the null. Phipson and Smyth (2010) claim that adding the identity can actually make the permutation test exact, i.e. of level α exactly. They do not actually prove this claim, but refer to the Monte Carlo test, which is exact in some situations if the original observation is added (Phipson and Smyth, 2010, and references therein).

Referring to Monte Carlo is not sufficient, because despite being related, a Monte Carlo test is very different from a permutation test. Monte Carlo samples are draws from the null distribution. In the permutation context there is not necessarily a unique null distribution, and the random permutations of the data are instead drawn from a conditional null distribution, i.e. the permutation distribution. It is however not obvious under what assumptions the original observation can also be seen as a random draw from this distribution, as in Monte Carlo testing.

Thus the proof by Phipson and Smyth (2010) is incomplete. As a result it remains unclear what assumptions (e.g. a group structure) are essential for the validity of their tests. Moreover, outside of the specific permutation scenario that they consider, it remains unknown if an exact test could be obtained. For example, it is unclear from Phipson and Smyth that random sampling from balanced permutations would lead to anti-conservative tests.

In this paper we investigate fundamental properties of permutation tests with random permutations. Our main focus is on the level of tests. Other properties, e.g. consistency, do not generally hold but can be established for more specific scenarios (Lehmann and Romano, 2005; Pesarin, 2015; Pesarin and Salmaso, 2013) by using results presented here. We discuss various ways of sampling random permutations and obtaining exact tests. In the literature there are two approaches to proving validity of permutation tests: a conditioning-based approach (Pesarin, 2015) and a more direct approach (Hoeffding, 1952; Lehmann and Romano, 2005). We will give proofs with both approaches since they complement each other in explaining permutation tests.

The structure of the paper is as follows. In Section 2 we review known results on the level of a permutation test based on a fixed group of transformations. The concepts and definitions from Section 2 are important throughout the paper. Testing with random permutations is covered in Section 3. In Section 3.1 we discuss the difference between permutation and Monte Carlo tests. Calculation of p-values when permutations are randomly drawn, is explained in Sections 3.2-3.4. In Section 3.4 we discuss novel results on exact tests with random permutations, which we prove in Section 3.5.

2 Testing with fixed transformations

This section deals with tests that employ the full group of transformations. Permutation tests are covered in Sections 2.1 and 2.2 and p-values are discussed in Section 2.3.

2.1 Exactness of the basic permutation test

We will emphasize the importance of the group structure of the set of transformations used. Here, and in the rest of the paper, we use the word 'group' in the strict algebraic sense, rather than loosely in the meaning of 'set' as is common in the statistics literature. We use the term 'permutation test' even if other transformations than permutations are considered.

Throughout this paper, we consider the following setting. Let X be data taking values in a sample space $\mathcal{X}$. Let G be a finite set of transformations $g : \mathcal{X} \to \mathcal{X}$, such that G is a group with respect to the operation of composition of transformations. The latter assumption is made to guarantee that

$Gg = G$ for all $g \in G$, which is used in all proofs. Considering a general group of transformations rather than simply the set of all permutations is useful, since in many practical situations the group consists of e.g. rotations (Langsrud, 2005; Solari et al., 2014) or maps that multiply part of the data by -1 (Pesarin and Salmaso (2010), pp. 54 and 168). We write $g(X)$ as gX. Consider any test statistic $T : \mathcal{X} \to \mathbb{R}$. We are concerned with testing any null hypothesis H_0 which implies that the joint distribution of the test statistics $T(gX)$, $g \in G$, is invariant under all transformations in G of X. That is, writing $G = \{a_1, ..., a_{\#G}\}$, under H_0

$$\big(T(a_1X), ..., T(a_{\#G}X)\big) \stackrel{d}{=} \big(T(a_1gX), ..., T(a_{\#G}gX)\big)$$

for all $g \in G$. Note that the latter holds in particular when for all $g \in G$

$$X \stackrel{d}{=} gX.$$

Composite null hypotheses are usually not of the form H_0, but for specific scenarios, properties of tests of such hypotheses can be established using results in this paper.

The most basic permutation test rejects H_0 when $T(X) > T^{(k)}(X)$, where

$$T^{(1)}(X) \leq ... \leq T^{(\#G)}(X)$$

are the sorted test statistics $T(gX)$, $g \in G$, and $k = \lceil (1-\alpha)\#G \rceil$ with $\alpha \in [0, 1)$. As is known and stated in the following theorem, this test has level at most α.

Theorem 1. *Under* H_0, $\mathbb{P}\big\{T(X) > T^{(k)}(X)\big\} \leq \alpha$.

We give two proofs of Theorem 1 in Section 2.2. The test of Theorem 1 is not always exact. When the data are discrete then the basic permutation test is often slightly conservative, due to a non-zero probability of tied values in X. However, under the following condition the test is exact for certain values of α.

Condition 1. There is a partition $\{G_1, ..., G_m\}$ of G with $id \in G_1$ and $\#G_1 = ... = \#G_m$, such that the following holds. Under H_0 with probability 1 for all $g, g' \in G$, $T(gX) = T(g'X)$ if and only if g and g' are in the same set G_i.

Proposition 1. *Under Condition 1, the test of Theorem 1 is exact for all* $\alpha \in \{0, 1/m, ..., (m-1)/m\}$.

The proof of this result is analogous to the proof of Theorem 1. Note that in Condition 1 each set G_i is an equivalence class of $\#G/m$ permutations that give the same test statistic. Such classes often exist. A simple application where Proposition 1 applies is the following. Consider a randomized trial where $X \in \mathbb{R}^{2n}$ and the test statistic is

$$T(X) = \sum_{i=1}^{n} X_i - \sum_{i=2n}^{n} X_i, \tag{1}$$

where $X_1, ..., X_n$ are cases and $X_{n+1}, ..., X_{2n}$ are controls and all X_i are independent and identically distributed under H_0. Let

$$m = \binom{2n}{n}.$$

If the observations are continuous then the set of α for which the test is exact is $\{0, 1/m, ..., (m-1)/m\}$, reflecting the fact that there are m equivalence classes of size $n!n!$ of permutations that always give the same test statistic.

When the basic permutation test of Theorem 1 is conservative due to e.g. tied values, it can be made exact by rejecting with a suitable probability a in the boundary case that $T(X) = T^{(k)}$, as Hoeffding (1952) shows. Here

$$a = a(X) = \frac{\alpha \#G - M^+(X)}{M^0(X)}, \tag{2}$$

where

$$M^+(X) := \#\{g \in G : T(gX) > T^{(k)}(X)\},$$
$$M^0(X) := \#\{g \in G : T(gX) = T^{(k)}(X)\}.$$

When G is not a group, the test can be highly anti-conservative or conservative. An example of a set of permutations that leads to an anti-conservative test, are balanced permutations. The set of balanced permutations is a subset of the set of all permutations, but is not a subgroup. These permutations have been used in various papers since they can have an intuitive appeal. They are discussed in Southworth et al. (2009), who warn against their use. The fact that permutations have been used incorrectly illustrates the importance of understanding and communicating very precisely which properties a permutation method should and should not have. In particular more emphasis should be put on the assumption of a group structure.

2.2 Proofs

We now turn to proving Theorem 1. We will give two proofs: a conditioning-based approach and an approach without conditioning. Both approaches are more or less known. We discuss both since each has distinct advantages.

The first approach has the advantage that it can be seen as more intuitive. The second approach on the other hand is more concise in proving the results in this paper. This is the case both for fixed permutations and random permutations. Another advantage of the second approach is that it leads to the discovery of a new random permutation sampling scheme in Section 3.4.

We now give the conditioning-based proof of Theorem 1. The proof is similar to that in Pesarin (2015) but the setting is more general. For each $x \in \mathcal{X}$, define O_x to be the orbit of x, which is the set $\{gx : g \in G\} \subseteq \mathcal{X}$.

Proof. Let $A = \{x \in \mathcal{X} : T(x) > T^{(k)}(x)\}$ be the set of elements of the sample space that lead to rejection. Suppose H_0 holds. By the group structure, $Gg = G$ for all $g \in G$. Consequently, $T^{(k)}(gX) = T^{(k)}(X)$ for all $g \in G$. Thus $\#\{g \in G : gX \in A\} =$

$$\#\{g \in G : T(gX) > T^{(k)}(gX)\} = \#\{g \in G : T(gX) > T^{(k)}(X)\} \leq \alpha \#G.$$

Endow the space of orbits with the σ-algebra that it naturally inherits from the σ-algebra on the sample space $\mathcal{X}$. Analogously to the proof of Theorem 15.2.2 in Lehmann and Romano (2005), we obtain

$$\mathbb{P}(X \in A \mid O_X) = \frac{1}{\#G} \#\{g \in G : gX \in A\}.$$

By the argument above, this is bounded by α. Hence

$$\mathbb{P}(X \in A) = E\big\{pr(X \in A \mid O_X)\big\} \leq \alpha$$

as was to be shown. □

We now state the algebraic proof. A similar proof can be found in Hoeffding (1952) and Lehmann and Romano (2005) (p. 634).

Proof. By the group structure, $Gg = G$ for all $g \in G$. Hence $T^{(k)}(gX) = T^{(k)}(X)$ for all $g \in G$. Let h have the uniform distribution on G. Then under

H_0, the rejection probability is

$$\begin{aligned}&\mathbb{P}\big\{T(X) > T^{(k)}(X)\big\} = \\ &\mathbb{P}\big\{T(hX) > T^{(k)}(hX)\big\} = \\ &\mathbb{P}\big\{T(hX) > T^{(k)}(X)\big\} \leq \alpha.\end{aligned}$$

The first equality follows from the null hypothesis and the second equality holds since $T^{(k)}(X) = T^{(k)}(hX)$. □

2.3 Permutation p-values

'Permutation p-value' is shorthand for a p-value based on permutations of the data. Here we will explain the meaning of such p-values and their calculation, when the set of permutations is fixed. Calculation of p-values based on random permutations is the subject of Section 3.2 and the end of Section 3.4. Estimation of p-values is covered in Section 3.3.

To understand the meaning of permutation p-values, it is essential to note that there is often no unique null distribution of $T(X)$, since in the permutation context H_0 often does not specify a unique null distribution of the data. Correspondingly, the permutation test is not a test that rejects H_0 when the test statistic $T(X)$ exceeds the $(1-\alpha)$-quantile of its null distribution.

When a test statistic t is a non-random function of the data and has a unique distribution under a hypothesis H, we define the 'p-value in the strict sense' as $\mathbb{P}_H(t \geq t_{obs})$, where t_{obs} is the observed value of t. A p-value in the strict sense always exists for some t, e.g. for t a constant function. For more meaningful test statistics however, such a p-value is not always defined. Since $T(X)$ does not always have a unique null distribution, no p-value in the strict sense can be based on it. However, under Condition 1 the statistic

$$D = \#\big\{g \in G : T(gX) \geq T(X)\big\}$$

does have a unique null distribution. Thus a p-value in the strict sense based on $-D$ is then defined. Denoting by d the observed value of D, we have under H_0

$$\mathbb{P}(-D \geq -d) = \mathbb{P}(D \leq d) = \mathbb{P}\big\{T(X) > T^{(\#G-d)}(X)\big\} = \frac{d}{\#G}.$$

This is indeed what is usually considered to be the permutation p-value. This equality holds under Condition 1. In other cases, such as when the observations are discrete, the null hypothesis often does not specify a unique null distribution of D. Thus there is not always a p-value in the strict sense based on D.

When H_0 does not specify a unique null distribution of any sensible test statistic, as a resolution a 'worst-case' p-value could be defined. However sometimes better solutions are possible, e.g. the randomized p-value p' in Section 3.4. In general, a 'p-value in the weak sense' can be considered, which we define as any random variable p satisfying $\mathbb{P}(p \leq c) \leq c$ for all $c \in [0, 1]$ for every distribution under some null hypothesis. For H_0, $D/\#G$ is always a p-value in the weak sense.

3 Testing with random transformations

In Section 3 we extend the results of the previous section to tests based on random transformations. Since permutation testing with random permutations is often confused with Monte Carlo testing, in Section 3.1 the differences between the two are made explicit. In Sections 3.2 and 3.3 calculation and estimation of p-values is discussed. Also the potential danger of not including the identity permutation is explained. Our results on exact tests are in Section 3.4 and their proofs in Section 3.5.

3.1 Comparison of Monte Carlo and permutation tests

In a basic Monte Carlo experiment, the null hypothesis H_0 is that X follows a specific distribution. A Monte Carlo test is used when there is no analytical expression for the $(1-\alpha)$-quantile of the null distribution of $T(X)$, such that the observed value of $T(X)$ cannot simply be compared to this quantile. To test H_0, independent realizations $X_2, ..., X_w$ are drawn from the null distribution of X. Assume that $T(X), T(X_2), ..., T(X_w)$ are continuous. Writing $X_1 = X$, let

$$B' = \#\{1 \leq j \leq w : T(X_j) \geq T(X)\}$$

and let b' denote its observed value. It is easily verified that under H_0, B' has the uniform distribution on $\{1, ..., w\}$.

The Monte Carlo test rejects H_0 when $T(X) > T^{(k')}$, where $k' = \lceil (1 - \alpha)w \rceil$ and $T^{(1)} \leq ... \leq T^{(w)}$ are the sorted test statistics $T(X_j)$, $1 \leq j \leq w$.

Note that $T^{(k')}$ is not the exact $(1-\alpha)$-quantile of the null distribution of $T(X)$. Instead, the principle behind the Monte Carlo test is the following. The null distribution of B' is known. So rather than comparing $T(X)$ to its null distribution, the test compares the statistic B' to its null distribution. That is, it rejects H_0 if and only if $-B'$ exceeds the $(1-\alpha)$-quantile of its null distribution, i.e. if and only if the Monte Carlo p-value

$$\mathbb{P}_{H_0}(B' \leq b') = b'/w,$$

where b' is the observed value of B', is at most α. This holds if and only if $T(X) > T^{(k')}$.

In the context of a permutation test based on random permutations, the validity of the test is not as clear as for Monte Carlo, since the random permutations of the data are not simply independent draws from the null distribution. Let $g_2, ..., g_w$ be random permutations from G. One permutation, g_1, is fixed to be $id \in G$, reflecting the original observation. Then, similarly to a Monte Carlo test, the permutation test rejects H_0 if and only if $T(X) > T^{(k')}(X)$, where now $T^{(1)} \leq ... \leq T^{(w)}$ are the sorted test statistics $T(g_j X)$, $1 \leq j \leq w$. Note that contrary to the Monte Carlo sample $X_1, ..., X_w$, the permutations $g_1 X, ..., g_w X$ are not independent under H_0.

This means that the essential property of the Monte Carlo sample which guarantees validity of the Monte Carlo test, does not hold for the random permutations of the data. Thus the validity of a test based on random permutations cannot be argued for by simply referring to Monte Carlo tests. To prove the validity of the test based on random permutations, we must use that $g_1 X, ..., g_w X$ are independent and identically distributed conditionally on the orbit O_X. It is however not clear when this is true: it is not obvious what properties G should have in order that $g_1 X = X$ can be seen as a random draw from O_X conditionally on O_X. It will be seen later that it suffices that G is a group.

Despite the fundamental difference, as suggested there is an analogy between tests with random permutations and Monte Carlo tests: the former can be said to be 'conditional Monte Carlo tests'. The test based on random permutations has level at most α if one permutation, e.g. g_1, is fixed to be the identity map, reflecting the original observation. As explained a similar result is known for Monte Carlo tests.

3.2 Calculating p-values based on random transformations

In practice it is often computationally infeasible to calculate the permutation p-value based on the *whole* permutation group, $D/\#G$. We distinguish two approaches in the literature for dealing with this. In both approaches, random permutations are used. The first approach is *calculating* (rather than estimating) a p-value based on the random permutations. This is discussed here and in Section 3.4. The second approach is *estimating* the p-value $D/\#G$, which is discussed in Section 3.3.

Phipson and Smyth (2010) give formulas for p-values, when permutations are randomly drawn. Under Condition 1, these formulas are valid if permutations are suitably sampled. Write

$$B = \#\big\{1 \leq j \leq w : T(\tilde{g}_j X) \geq T(X)\big\}, \tag{3}$$

where $\tilde{g}_1, ..., \tilde{g}_w$ are random permutations with distribution to be specified. Let b be the observed value of B. Under Condition 1 and suitable sampling, Phipson and Smyth's p-values are exactly equal to $\mathbb{P}_{H_0}(-B \geq -b)$. These are not p-values in the strict sense as defined in Section 2.3, since they depend on randomly drawn permutations. Arguably however, they are as close as possible to being p-values in the strict sense, given that random permutations are used. Under Condition 1, if $\tilde{g}_1, ..., \tilde{g}_w$ are drawn such that they are from distinct elements G_i of the partition and not from G_1, the p-value $\mathbb{P}_{H_0}(-B \geq -b)$ is exactly

$$\frac{b+1}{w+1}.$$

The validity of this formula follows from Proposition 2. For the case that permutations are drawn with replacement, where $\tilde{g}_1, ..., \tilde{g}_w$ are independent and uniform on G, Phipson and Smyth also provide a formula for $\mathbb{P}_{H_0}(-B \geq -b)$, under Condition 1.

Phipson and Smyth's formula for the p-value in the case of drawing without replacement implies that B/w is the p-value if one of the permutations is fixed to be the identity map, i.e. if $\tilde{g}_1$ is taken to be the identity. It follows that the permutation test based on random permutations becomes exact for certain α if the identity is added. Note that this only holds if Condition 1 is satisfied and all permutations are from distinct equivalence classes G_i.

The formula for the p-value that Phipson and Smyth provide for the case of drawing with replacement is somewhat involved but can be useful in practice. The proofs underlying their formulas however are incomplete. Indeed, it does not become clear why the structure of the set of all permutations is such that their claims hold, and whether other sets of permutations or other transformations would also lead to valid p-values. Lehmann and Romano (2005) (p.636) remark without proof that if G is a group, then under H_0 the p-value $(B+1)/(w+1)$ is always stochastically larger than uniform on $[0,1]$, but they state no other properties. In Section 3.4 and Section 3.5 we provide a theoretical foundation for tests and p-values based on random transformations.

3.3 Estimated p-values

We now discuss the use of random permutations for estimating the p-value $D/\#G$. The random permutations $\tilde{g}_1, ..., \tilde{g}_w$ are typically all taken to be uniform on G and can be drawn with or without replacement. The estimate of p is often taken to be $\hat{p} = B/w$, with B as defined above. This is an unbiased estimate of p, i.e. $E\hat{p} = p$, and usually $\lim_{w\to\infty} \hat{p} = p$.

The more conservative estimate $\tilde{p} = (B+1)/(w+1)$ could of course also be used, but is not unbiased. On the other hand, the potential advantage of $\tilde{p}$ is that it satisfies $\mathbb{P}(\tilde{p} \leq c) \leq c$ for all $c \in [0,1]$ under H_0, as follows from Theorem 2.

Using the unbiased estimate $\hat{p} = B/w$ can be very dangerous, as Phipson and Smyth (2010) thoroughly explain. The reason is that $\hat{p}$ is almost never stochasticaly larger than the uniform distribution on $[0,1]$ under H_0. This is immediately clear from the fact that $\hat{p}$ usually has a strictly positive probability of being zero. Consequently, if H_0 is rejected if $\hat{p} \leq c$ for some cut-off c, then the type-I error rate can be larger than c. Often this difference will be small for large w. However, when c is itself small due to e.g. Bonferroni's multiple testing correction, then $\mathbb{P}(\hat{p} \leq c)$ can become many times larger than c under H_0. This is because this probability does not converge to zero as $c \downarrow 0$ for fixed w. Thus, as Phipson and Smyth (2010) note, using $\hat{p}$ in combination with e.g. Bonferroni can lead to completely faulty inference. Appreciable anti-conservativeness also occurs if very few (e.g. 25–100) random permutations are used (as in e.g. Byrne et al. (2013) and Schimanski et al. (2013)).

3.4 Tests based on random permutations

Here we discuss the main results of the paper. Proofs are in Section 3.5. Theorem 2 states that the permutation test with random permutations has level at most α if the identity map is added. This has been remarked, but without proof. We first define the vector of random transformations.

Definition 1. Let G' be the vector $(id, g_2, ..., g_w)$, where id is the identity in G and $g_2, ..., g_w$ are random elements from G. Write $g_1 = id$. The transformations can be drawn either with or without replacement: the statements in this paper hold for both cases. If we draw $g_2, ...g_w$ *without* replacement, then we take them to be uniformly distributed on $G \setminus \{id\}$, otherwise uniform on G.

Theorem 2. *Let G' be as in Definition 1. Let $T^{(1)}(X, G') \leq ... \leq T^{(w)}(X, G')$ be the ordered test statistics $T(g_j X)$, $1 \leq j \leq w$. Let $\alpha \in [0, 1)$ and recall that $k' = \lceil (1 - \alpha)w \rceil$.*

Let H_0 be as in Section 2.1. Reject H_0 when $T(X, G') > T^{(k')}(X, G')$. Then the rejection probability under H_0 is at most α.

Theorem 2 implies that $(B + 1)/(w + 1)$ is always a p-value in the weak sense if $\tilde{g}_1, ..., \tilde{g}_w$ are all randomly drawn with replacement from G or without replacement from $G \setminus \{g_1\}$. In specific scenarios, Theorem 2 can be extended to certain composite null hypotheses. Proposition 2 states that under Condition 1 and suitable sampling, the test with random permutations is exact. The formula in Section 3.2 for the p-value under sampling without replacement is equivalent to the last part of this result.

Proposition 2. *Suppose Condition 1 holds. Let $h_1 \in G_1, ..., h_m \in G_m$. Then the result of Theorem 2 still holds if $g_2, ..., g_w$ are drawn with replacement from $\{h_1, ..., h_m\}$ or without replacement from $\{h_2, ..., h_m\}$. Moreover, in the latter case, the test of Theorem 2 is exact for all $\alpha \in \{0/w, 1/w, ..., (w - 1)/w\}$.*

Using this result it can also be shown that specific tests with random permutations are unbiased. The test of Theorem 2 can be slightly conservative if α is not chosen suitably or due to the possibility of ties. Recall that the same holds for the basic permutation test that uses all transformations in G. The adaptation by Hoeffding at (2) then guarantees exactness. The following is a generalization of Hoeffding's result to random transformations.

Proposition 3. *Consider the setting of Theorem 2. Let*

$$a = a(X, G') = \frac{w\alpha - M^+(X, G')}{M^0(X, G')}, \tag{4}$$

where

$$M^+(X, G') := \#\{1 \le j \le w : T(g_j X) > T^{(k')}(X, G')\},$$
$$M^0(X, G') := \#\{1 \le j \le w : T(g_j X) = T^{(k')}(X, G')\}.$$

Reject if $T(X) > T^{(k')}(X, G')$ and reject with probability a if $T(X) = T^{(k')}(X, G')$. Then the rejection probability is exactly α under H_0.

The proof of this result is in Section 3.5. The test of Proposition 3 entails a randomized decision: in case $T(X) = T^{(k')}$, the test randomly rejects with probability a. This is in itself not objectionable, since the test is randomized anyway due to the random transformations. Note that in the situation of Proposition 2 under drawing without replacement the test is already exact, such that Proposition 3 is not needed to obtain an exact test.

In Theorem 2, the requirement of using the whole group is replaced by suitable random sampling from the group. Alternative sampling schemes are possible. For example, let $G^* \subseteq G$ be any finite subset of G, where we now allow G to be an infinite group as well. Write $k^* = \lceil (1-\alpha)\#G^* \rceil$. Let h be uniformly distributed on G^* and independent. Let H_0 be as in Theorem 2 but reject H_0 if and only if

$$T(X) > T^{(k^*)}(X, G^* h^{-1}),$$

i.e. if $T(X)$ exceeds the $(1-\alpha)$-quantile of the values $T(gh^{-1})$, $g \in G^*$. This is a randomized rejection rule, since it depends on h, which is randomly drawn each time the test is executed. The rejection probability is at most α, which follows from an argument analogous to the last four steps of the second proof of Theorem 2. Note that if G^* is a group itself, then $G^* h^{-1} = G^*$ and this test becomes non-random, coinciding with the basic permutation test. Thus it is a generalization thereof. This result allows using a permutation test when G is an infinite group of transformations, from which it may not be obvious how to sample uniformly. One simply uses any finite subset G^* of the infinite group.

Corresponding to the randomized test of Proposition 3, a randomized p-value can be defined as follows. The advantage of this p-value is that

it is always uniform on $[0,1]$ under H_0 without requirement of additional assumptions, and it is is easy to calculate. Consider the randomized test of Proposition 3. Suppose without loss of generality that when $T(X) = T^{(k')}$, the test rejects if and only if $a > u$, where u is uniform on $[0,1]$ and independent. Define the randomized p-value by

$$p' = \frac{\#\{1 \leq j \leq w : T(g_j X) > T(X)\}}{w} + u\frac{\#\{1 \leq j \leq w : T(g_j X) = T(X)\}}{w}.$$

This p-value has the property that $p' \leq \alpha$ if and only if the randomized test rejects. This implies in particular that p' is exactly uniform on $[0,1]$ under H_0. It is a p-value in the weak sense as defined in Section 2.3. As for the test of Proposition 3, the fact that p' is randomized is in itself not objectionable, since it is randomized anyway due to the random transformations.

A simple upper bound to p' is

$$\frac{\#\{1 \leq j \leq w : T(g_j X) \geq T(X)\}}{w},$$

another p-value in the weak sense, which is equivalent with $(B+1)/(w+1)$ when e.g. $\tilde{g}_1, ..., \tilde{g}_w$ are all independent uniform draws from G. It is not exactly uniform on $[0,1]$ under H_0. However, when w is large and there are few ties among the test statistics, it tends to closely approximate p', so that it may be used for simplicity.

3.5 Proofs

We provide two proofs of Theorem 2. The first of these is a conditioning-based approach.

Proof. We prove the result for the case of drawing with replacement. The proof for drawing without replacement is analogous. Note that (X, G') takes values in $\mathcal{X} \times \{id\} \times G^{w-1}$. Let $A \subset \mathcal{X} \times \{id\} \times G^{w-1}$ be such that the test rejects if and only if $(X, G') \in A$.

Endow the space of orbits with the σ-algebra that it naturally inherits from the σ-algebra on $\mathcal{X}$. Suppose H_0 holds. Assume that almost surely O_X contains $\#G$ distinct elements. In case not, the proof is analogous. Analogously to the proof of Theorem 15.2.2 in Lehmann and Romano (2005), we obtain

$$\mathbb{P}\big\{(X, G') \in A \mid O_X\big\} = \frac{\#\big(O_X \times \{id\} \times G^{w-1}\big) \cap A}{\#O_X \times \{id\} \times G^{w-1}}. \tag{5}$$

We will now argue that this is at most α. Fix X. Let $\tilde{X}$ have the uniform distribution on O_X. It follows from the group structure of G that the entries of $G'\tilde{X}$ are just independent uniform draws from O_X. Thus from the Monte Carlo testing principle it follows that $\mathbb{P}\{(\tilde{X}, G') \in A\} \leq \alpha$. Since $(\tilde{X}, G')$ was uniformly distributed on $O_X \times \{id\} \times G^{w-1}$, it follows that (5) is at most α. We conclude that

$$\mathbb{P}\big\{(X, G') \in A\big\} = E\Big[\mathbb{P}\big\{(X, G') \in A \mid O_X\big\}\Big] \leq \alpha,$$

as was to be shown. □

We now give a proof of Theorem 2 without conditioning.

Proof. From the group structure of G, it follows that for all $1 \leq j \leq w$, $G'g_j^{-1}$ and G' have the same distribution, if we disregard the order of the elements. Let j have the uniform distribution on $\{1, ..., w\}$ and write $h = g_j$. Under H_0,

$$\begin{aligned}
&\mathbb{P}\big\{T(X) > T^{(k')}(X, G')\big\} = \\
&\mathbb{P}\big\{T(X) > T^{(k')}(X, G'h^{-1})\big\} = \\
&\mathbb{P}\big\{T(hX) > T^{(k')}(hX, G'h^{-1})\big\}.
\end{aligned}$$

Since $(G'h^{-1})(hX) = G'(h^{-1}hX)$, the above equals

$$\begin{aligned}
&\mathbb{P}\big\{T(hX) > T^{(k')}(h^{-1}hX, G')\big\} = \\
&\mathbb{P}\big\{T(hX) > T^{(k')}(X, G')\big\} \leq \alpha
\end{aligned}$$

as was to be shown. □

We now prove Proposition 2.

Proof. We consider the case that $g_2, ..., g_w$ are drawn without replacement from $\{h_2, ..., h_m\}$ and show that the test is exact for $\alpha \in \{0/w, ..., (w-1)/w\}$. Write $G' = (g_1, ..., g_w)$. Let h have the uniform distribution on $\{g_1, ..., g_w\}$. For each $g \in G$ let $i(g) \in \{1, ..., m\}$ be such that $g \in G_{i(g)}$. Suppose H_0 holds. From the group structure of G it follows that the sets $\big\{i(g_1), ..., i(g_w)\big\}$ and $\big\{i(g_1h^{-1}), ..., i(g_wh^{-1})\big\}$ have the same distribution. Consequently

$$\begin{aligned}
&\mathbb{P}\big\{T(X) > T^{(k')}(X, G')\big\} = \\
&\mathbb{P}\big\{T(X) > T^{(k')}(X, G'h^{-1})\big\}.
\end{aligned}$$

As in the above proof of Theorem 2 we find that this equals $\mathbb{P}\{T(hX) > T^{(k')}(X, G')\}$.

Since $\alpha \in \{0/w, ..., (w-1)/w\}$ and $T(g_1X), ..., T(g_wX)$ are distinct, it holds with probability one that

$$\#\{1 \leq j \leq w : T(g_jX) > T^{(k')}\} = \alpha w.$$

Since h is uniform, it follows that $\mathbb{P}\{T(hX) > T^{(k')}(X, G')\} = \alpha$. □

We now prove Proposition 3.

Proof. Note that

$$\mathbb{P}(\text{reject}) = E\{\mathbb{1}_{\{T(X)>T^{(k')}(X,G')\}} + a(X, G')\mathbb{1}_{\{T(X)=T^{(k')}(X,G')\}}\}.$$

Write $M^+ = M^+(X, G')$ and $M^0 = M^0(X, G')$. Analogously to the first four steps of the second proof of Theorem 2, it follows for h as defined there that the above equals

$$\begin{aligned}
&E\{\mathbb{1}_{\{T(hX)>T^{(k')}(X,G')\}} + a(X, G')\mathbb{1}_{\{T(hX)=T^{(k')}(X,G')\}}\} = \\
&E\{\mathbb{1}_{\{T(hX)>T^{(k')}(X,G')\}}\} + E\{a(X, G')\mathbb{1}_{\{T(hX)=T^{(k')}(X,G')\}}\} = \\
&E\{M^+w^{-1}\} + E\Big[E\Big\{\frac{w\alpha - M^+}{M^0}\mathbb{1}_{\{T(hX)=T^{(k')}(X,G')\}} \mid M^+, M^0\Big\}\Big] = \\
&E\{M^+w^{-1}\} + E\Big[\frac{w\alpha - M^+}{M^0}E\Big\{\mathbb{1}_{\{T(hX)=T^{(k')}(X,G')\}} \mid M^+, M^0\Big\}\Big] = \\
&E\{M^+w^{-1}\} + E\Big[\frac{w\alpha - M^+}{M^0}M^0w^{-1}\Big] = \alpha,
\end{aligned}$$

as was to be shown. □

Discussion

This paper focuses on proving properties of tests with random permutations under relatively general conditions. Properties such as unbiasedness of tests of composite null hypotheses and consistency do not hold in general but may be proven for more specific scenarios. For fixed permutations, there are many results regarding such properties (Hoeffding, 1952; Lehmann and Romano, 2005; Pesarin and Salmaso, 2010, 2013) which may be extended to random permutations.

Aside from the permutation test, there are many multiple testing methods which employ permutations. Many of these are precisely based on the principle behind the permutation test (Westfall and Young, 1993; Tusher et al., 2001; Meinshausen and Bühlmann, 2005; Meinshausen, 2006). This paper can provide better insight into these procedures, when random permutations are used. Many results in this paper have extensions to these methods.

Acknowledgements

We thank Aldo Solari and Vincent van der Noort for their valuable suggestions.

References

Byrne, E., Carrillo-Roa, T., Henders, A., Bowdler, L., McRae, A., Heath, A., Martin, N., Montgomery, G., Krause, L., and Wray, N. Monozygotic twins affected with major depressive disorder have greater variance in methylation than their unaffected co-twin. *Translational psychiatry*, 3(6):e269, 2013.

Dwass, M. Modified randomization tests for nonparametric hypotheses. *The Annals of Mathematical Statistics*, 28:181–187, 1957.

Ernst, M. D. et al. Permutation methods: a basis for exact inference. *Statistical Science*, 19(4):676–685, 2004.

Fisher, R. A. "the coefficient of racial likeness" and the future of craniometry. *Journal of the Anthropological Institute of Great Britain and Ireland*, 66: 57–63, 1936.

Ge, Y., Dudoit, S., and Speed, T. P. Resampling-based multiple testing for microarray data analysis. *Test*, 12(1):1–77, 2003.

Goeman, J. J. and Solari, A. The sequential rejection principle of familywise error control. *The Annals of Statistics*, 38:3782–3810, 2010.

Hoeffding, W. The large-sample power of tests based on permutations of observations. *The Annals of Mathematical Statistics*, 23:169–192, 1952.

Langsrud, Ø. Rotation tests. *Statistics and computing*, 15(1):53–60, 2005.

Lehmann, E. L. and Romano, J. P. *Testing statistical hypotheses*. Springer Science & Business Media, 2005.

Meinshausen, N. False discovery control for multiple tests of association under general dependence. *Scandinavian Journal of Statistics*, 33(2):227–237, 2006.

Meinshausen, N. and Bühlmann, P. Lower bounds for the number of false null hypotheses for multiple testing of associations under general dependence structures. *Biometrika*, 92(4):893–907, 2005.

Pesarin, F. Some elementary theory of permutation tests. *Communications in Statistics-Theory and Methods*, 44(22):4880–4892, 2015.

Pesarin, F. and Salmaso, L. On the weak consistency of permutation tests. *Communications in Statistics-Simulation and Computation*, 42(6):1368–1379, 2013.

Pesarin, F. and Salmaso, L. *Permutation tests for complex data: theory, applications and software*. John Wiley & Sons, 2010.

Phipson, B. and Smyth, G. K. Permutation p-values should never be zero: calculating exact p-values when permutations are randomly drawn. *Statistical applications in genetics and molecular biology*, 9(1):39, 2010.

Schimanski, L. A., Lipa, P., and Barnes, C. A. Tracking the course of hippocampal representations during learning: when is the map required? *The Journal of Neuroscience*, 33(7):3094–3106, 2013.

Solari, A., Finos, L., and Goeman, J. J. Rotation-based multiple testing in the multivariate linear model. *Biometrics*, 70(4):954–961, 2014.

Southworth, L. K., Kim, S. K., and Owen, A. B. Properties of balanced permutations. *Journal of Computational Biology*, 16(4):625–638, 2009.

Tusher, V. G., Tibshirani, R., and Chu, G. Significance analysis of microarrays applied to the ionizing radiation response. *Proceedings of the National Academy of Sciences*, 98(9):5116–5121, 2001.

Westfall, P. H. and Young, S. S. *Resampling-based multiple testing: Examples and methods for p-value adjustment*. John Wiley & Sons, 1993.